

AUTONOMY ENABLES NEW SCIENCE MISSIONS

Richard J. Doyle
Victoria Gor
Guy K. Man
Paul E. Stolorz
Jet Propulsion Laboratory
Pasadena, CA 91109

Clark Chapman
William J. Merline
Alan Stern
Southwest Research Institute
Boulder, CO 30303

Abstract

The challenge of space flight in **NASA's** future **is** to enable smaller, more frequent and intensive space exploration at much lower total cost without substantially decreasing mission reliability, capability, or the scientific return on investment. The most effective way to achieve this goal is to build intelligent capabilities into the spacecraft themselves. Our technological vision for meeting the challenge of returning quality science through limited communication bandwidth will actually put scientists in a more direct link with the spacecraft than they have enjoyed to date. Technologies such as pattern recognition and machine learning can place a part of the scientist's awareness onboard the spacecraft to prioritize downlink of to autonomously trigger time-critical follow-up observations---particularly important in flyby missions---without ground interaction. Onboard knowledge discovery methods can be used to include candidate discoveries in each downlink for scientists' scrutiny. Such capabilities will allow scientists to quickly reprioritize missions in a much more intimate and efficient manner than is possible today. Ultimately, new classes of exploration missions will be enabled.

INTRODUCTION

Highly Autonomous Platforms

We envision that intelligent autonomous space systems will evolve and deploy in major phases. The first phase involves automation of the basic engineering and mission accomplishment functions of the space platform. The relevant capabilities include mission planning and resource management, health management and fault protection, and guidance, navigation and control. Stated differently, these autonomous capabilities will make the space platform *self-directing*, *self-preserving* and *self-mobilizing*. Some of the relevant technologies include **Artificial Intelligence (AI)** based planning & scheduling, operations research, decision theory, model-based reasoning, intelligent agents, spatial reasoning and neural and other specialized technologies. By 2000, we expect that NASA spacecraft will have achieved an onboard automated closed loop control at a basic level among: planning activities to achieve mission goals, maneuvering and pointing to execute those activities, and detecting and resolving faults to continue the mission without requiring ground support. At this point, mission accomplishment is becoming largely autonomous, and dramatic cost savings is seen in the form of reduced, shared ground staffing which responds on demand to the spacecraft's beacon-based requests for interaction. Also in this phase, the first elements of onboard science autonomy will be deployed, using technologies like trainable object recognizers. However, the decision-making capacity to determine how mission priorities should change and what new mission goals should be added in the light of intermediate results, discoveries and other events would still reside largely with scientists and other analysts on the ground.

Observing and Discovery Presence Onboard

Work on automating the spacecraft will continue into challenging areas like greater onboard adaptability in responding to events, closed-loop control for small body rendezvous and landing missions, and operation of the multiple free-flying elements of space-based telescopes and interferometers. In addition, in the next phase of autonomy development and insertion, a portion of the scientist's awareness, i.e., an observing and discovery presence, will move aggressively onboard. In other words, knowledge for discriminating and

determining what information is important would migrate to the space platform. The relevant capabilities include feature detection and tracking, object recognition, and exploratory data sampling. Some of the relevant technologies are pattern recognition, machine learning data mining and knowledge discovery. At this point, the space platform begins to become *self-educating*, and can respond to uncertainty within the mission context, a prerequisite for graduating beyond reconnaissance to interactive, *in situ* exploration. By 2005, we expect that a significant portion of the information routinely returned from platforms would not simply and strictly match features of stated prior interest, but would be deemed by the onboard software to be "interesting" and worthy of further examination by appropriate experts on the ground. At this point, limited communications bandwidth would then be utilized in an extremely efficient fashion, and "alerts" from various and far-flung platforms would be anticipated with great interest.

This paper will focus on concepts and preliminary work for the second phase of autonomy deployment as outlined above, for projecting the scientist's awareness and creating an observing and discovery presence onboard the spacecraft. In the remainder of the paper, we first describe ongoing work on spacecraft autonomy and how such automation serves science, then we describe our concepts for applying autonomy technology for onboard science. Next we describe two onboard science experiments in detail, and we conclude with our thoughts on the future prospects and payoff for autonomy in the service of science and exploration.

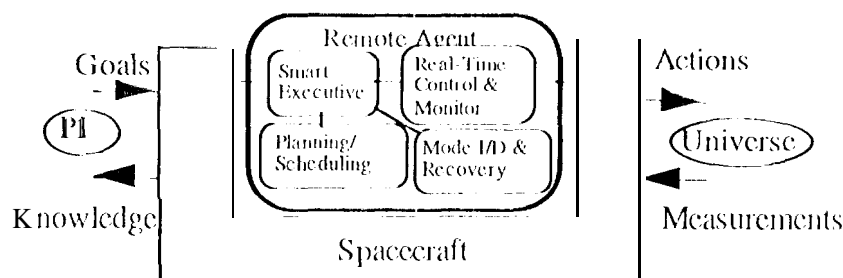
AUTONOMY FOR ENGINEERING FUNCTIONS

Automation of engineering functions is needed before automation of science functions because they provide the required infrastructure for onboard science. Without it, science autonomy would have limited success and benefits. Much of the initial focus for spacecraft autonomy has been on developing new software and systems concepts to automate engineering functions of the spacecraft: planning and scheduling, operations, guidance, navigation and control, fault protection and resource management. However, the ultimate objectives of NASA missions are science-driven objectives.

Immediate advances are being made in the development of a system concept and an architecture for autonomy software (Fesq 1996). This architecture covers the flight and ground system for all functions of the spacecraft. Since the focus of this paper is flight system autonomy, we will discuss the flight system portion of the architecture. It is important to look at autonomy from a system viewpoint because this approach gives the most gain in reducing operations cost and improving mission capability. The reason is because due to limited onboard resources, spacecraft functions are highly coupled and automating one part of the flight system with no consideration of the whole will not yield much savings. For example, retargeting the science camera for remote sensing has to be coordinated with attitude control, spacecraft communication, power distribution and data management, etc. All these functions have to be managed to provide conflict-free operations of the science camera. Automating attitude control alone is only accomplishing part of the job. This is precisely the reason why autonomy needs to be approached top-down using a system approach. Figure 1 is a block diagram of the spacecraft. The U-shaped block represents the spacecraft hardware and the middle square represents the flight software. A high degree of spacecraft autonomy is made possible by an a-run-time architecture consisting of four parts: planning/scheduling, smart executive, mode identification and reconfiguration, and real-time control. The real-time control segment is where the traditional real-time flight software such as attitude control resides and we will refer to the subsystem software as domain experts. The planning/scheduling, the smart executive and the mode identification and reconfiguration work closely together and are known collectively as the Remote Agent. This is the top application layer of the flight software managing the operations of the spacecraft functions and it is the core of our new autonomy architecture.

The Remote Agent's design is based on powerful and exciting technologies which endow the spacecraft with new behavior and capabilities. The Remote Agent carries explicit models of the operational behaviors of the spacecraft such as state models of the hardware, and resource requirements and operational constraints. The reasoning engines in the Remote Agent look for conflict free procedures to operate the spacecraft in real-time. The benefits of this approach is to reduce operations requirements while bestowing the spacecraft with multiple means to function in uncertain environments. Constraint-based planning and scheduling ensures achievement of long-term mission objectives and manages the allocation of system

resources. The smart executive performs robust, multi-threaded execution, to reliably execute planned sequences under conditions of uncertainty, to rapidly respond to unexpected events such as component failures, and to manage concurrent real-time activities. Mode identification and reconfiguration, based on model-based diagnosis, confirms successful plan execution and infers the health of all system components based on inherently limited sensor information. The three engines work together transforming the spacecraft from an open loop system (as seen by the ground operator) to a closed-loop system. The benefit of the closed-loop approach allows the spacecraft to go into environments that are more uncertain and to achieve our science goals more reliably in the face of such uncertainty.



Run-Time Architecture

Figure 1. The Remote Agent and Its Context

The Remote Agent also provides a scalable modular architecture for flight software. Figure 2 shows a general purpose superstructure. Specialized functions are delivered as custom, modular domain experts such as a science camera expert or an attitude control expert. These functions can be traditional spacecraft subsystem functions or they can be domain specific autonomous functions incorporating new technology such as optical navigation. The Remote Agent provides centralized autonomous functions supporting all specialized functions. This flexible and scalable architecture can serve a wide variety of systems, both in Space and on the ground.

AUTONOMY FOR ONBOARD SCIENCE

Automating the engineering functions of a spacecraft is already happening, and is tremendously challenging, but it is not an end in itself. Rather, the motivation for the development of spacecraft autonomy ultimately is to serve science. In the next sections, we discuss additional ideas for how autonomy technologies can automate some forms of science data processing onboard, and how such capabilities can interact with spacecraft autonomy capabilities to enhance science return from missions of exploration. We are proposing a new framework for autonomous evaluation of science data and observation planning onboard spacecraft. The future NASA mission set will feature smaller and more numerous spacecraft in an environment of highly constrained uplink and downlink communications. The proposed paradigm will strike a new, more ambitious balance among: direction of mission activities by scientists without the assistance of a ground sequencing team, robust capture and redirection in making discoveries at the target body, accommodation of the realities of limited communication links, and the return of quality science products from missions. Our concepts are aimed at pursuing autonomy for science as aggressively as autonomy for engineering, and within the same early time frame. We are developing a number of onboard science experiments in this vein.

The specific objectives of these initial experiments are as follows: (1) To demonstrate the ability to autonomously identify features and objects of known interest in onboard acquired data and to prioritize and/or edit downlink on the basis of reliably recognizing such features and objects. (2) To provide the basis for capturing transient science events through integration of autonomous onboard science data processing with autonomous onboard capabilities for retargeting and mission planning. (3) To provide the basis for scientists to efficiently redirect mission activities following scientific discoveries at the target body.

The intelligent systems technologies that will enable autonomy for science are the data mining technologies, including pattern recognition, machine learning and knowledge discovery techniques,

combined with other capabilities of an autonomous spacecraft, particularly onboard planning. We are building on previous work on ground-based automated science data processing, notably the SKICAT system for automating the generation of a comprehensive sky object catalog from Mt. Palomar observatory data (Djorgovski 1994, Fayyad 1993) and the JARTool system for automating the detection of volcanos in SAR data returned from the Magellan spacecraft at Venus (Burl 1994, Head 1991).

Benefits

Telemetry limitations place extreme constraints on the scope of scientific experiments possible for deep space missions; such constraints will become even more severe in the coming era of proliferation of deep space missions. Our aim is to demonstrate the ability of novel algorithms implemented on advanced flight computers to directly enhance the results achievable by scientific experiments onboard spacecraft. We plan to achieve this goal by implementing onboard data analysis algorithms that can

- 1) Rapidly sift through instrument data immediately upon collection,
- 2) Provide a massively condensed summary of the important information collected by the instrument(s), either to science DPs on the ground, or to an onboard planning executive, and
- 3) Enable adaptive control of observations based upon immediate data processing and analysis.

The goal is to generate information that fits within telemetry limitations, but that is nevertheless reliable and sufficient for the PI or an onboard planner to adaptively direct the spacecraft so that phenomena of special interest can be focused upon by spaceborne instruments. One of the benefits, therefore, the results of onboard analysis will be to achieve data contraction in downlink of several orders of magnitude.

Data contraction can take several forms. One obvious form is that of standard data compression algorithms, both lossless and lossy. The leverage available here should be pursued, but our main focus will be on more intelligent methods of data contraction. For example,

- 1) Data editing to transmit images to ground at high resolution, but only of those regions from an original image that are deemed to contain significant or unexpected scientific information.
- 2) Retargeting of the spacecraft to study important areas, following successful onboard feature detection.
- 3) Rapid downlink to Earth of potentially interesting target regions and phenomena selected by onboard analysis software.

Verification of discoveries by scientists on the ground can be followed by training and uploading of new recognizers to quickly extend the suite of autonomously detectable features onboard, a much more efficient and intimate mission redirection process than currently enjoyed by scientists.

Proposed Experiments

The two onboard science experiments which follow have been proposed in the context of the New Millennium Deep Space 1 (DS-1) mission, a technology validation mission which will visit a comet and an asteroid. Each experiment involves the autonomous recognition of certain features of scientific interest at cometary and asteroidal targets. The aim of the experiments is to demonstrate how such a capability can be used to prioritize downlink so that images that contain such features are reliably identified and downlinked first. The same capability, if integrated with an onboard planner and pointing capability, would be the basis of autonomously retargeting and collecting additional images of the features of interest.

SATELLITE SEARCH

Science Objectives

The possibility of satellites orbiting asteroids has been debated since the late 1970s, when anomalous observations of stars being occulted near asteroids were first reported (Van Flandern et al., 1979). Theoretical studies of how asteroids formed and how they break up in collisions have yielded expectations that satellites would be absent (Gehrels et al., 1987) to uncommon (Weidenschilling et al., 1989). The discovery by the Galileo spacecraft in 1993 of one---and only one---satellite orbiting around Ida, one of

only two asteroids observed close-up by spacecraft, has spurred new thinking about how the asteroids collisionally evolve. For instance, Dunda (1996) now believes that formation of *small* satellites may be fairly common in catastrophic collisions. It remains a fascinating question about whether swarms of satellites could exist, perhaps segregated into separate resonances, or whether chaotic orbits would cause multiple satellites to quickly coalesce into a single one, like Ida's satellite Dactyl.

One of the major objectives of any mission to an asteroid will be to set constraints on the number and sizes of any satellites. The size, orbital parameters, and composition of a satellite hold significant clues to the origin and age of the asteroid itself and have larger implications for understanding asteroids as a whole, and consequently their role in the evolution of our solar system.

Scientific Benefits

Except for those images taken far from the asteroid, there will be essentially no time, in a flyby mission, to transmit images to Earth, search for satellites, and send commands for retargeting and sequencing to the spacecraft. Even in a rendezvous mission, the volume of space to be searched is large and the observing geometry -- especially with respect to small, nearby satellites -- changes very rapidly. An efficient search for satellites therefore will depend substantially on development of the types of automation proposed here.

The resulting benefits of an automated detection process will be to provide the opportunity to obtain additional high-resolution images, to study cratering history, satellite shape, and surface geology. Spectral data can then be acquired to study satellite composition and its relationship to the asteroid, other small bodies, and our inventory of meteorites. Opportunities will be available to follow the satellite over sufficient duration and from various perspectives to determine the satellite's orbit and thereby determine the mass of the central asteroid. This would be in marked contrast to the serendipitous observations of Dactyl by Galileo, which could only provide constraints -- rather than definitive determinations -- on Dactyl's orbit and thus on Ida's mass and bulk density.

Flyby missions to specific asteroids and comets (unlike missions to planets, to which we are likely to return), are effectively one-time events. Decades or perhaps centuries will intervene before spacecraft will return to these bodies. It is essential to extract as much scientific information as possible when the opportunity is at hand. Orbital missions to planets, asteroids, or comets will also benefit by making search strategies for small satellites more efficient. Therefore, all of NASA's future deep-space missions should benefit from the development of the technology proposed here.

Problem Description

The essence of the problem is detecting a real object in the presence of similar-appearing features: background stars, detector defects, and cosmic ray hits. The problem is complicated by the spacecraft velocity and rapidly changing observing geometry. Additionally, the satellites themselves will be moving in their orbits.

The simple approach might seem to be to take two frames in rapid sequence and look for those objects that remain relatively fixed. As our previous experience on Galileo showed, that is not sufficient. Cosmic rays can and do hit, even in five successive frames, in ways that can be mistaken for a real object.

Only frames taken from the far-field will encompass the entire Hill sphere, the sphere of gravitational influence surrounding any body. An advantage is that objects will not appear to move quickly due to spacecraft motion, but a disadvantage is that only a large satellite would be detectable.

When the spacecraft is closer, the detectability size-limit gets smaller, but the volume of space searched by a single frame also decreases. Any potential objects will move quickly in the field, perhaps being visible in only one or a few frames taken in rapid sequence.

Technical Approach

The approach to detection of small satellites is a straightforward one. It requires many successive frames to track the motion of potential objects. Such frame sequences are taken routinely for other mission or scientific objectives, including optical navigation and color sequences. No special images should be required to test the feasibility of this technique.

The stars in the field are easy to detect, both by their identifiable patterns and by their common motion in the field (their motion will be defined by the spacecraft trajectory and any camera tracking motion that is employed).

Cosmic rays, in the simplest sense, will hit in different places in different frames. But one must be aware that potential objects will move from one frame to the next also. Depending on the trajectory of the spacecraft and the position of the object with respect to the spacecraft and the target body (asteroid or comet), and, to a lesser extent, the object's orbital motion, the path of the object on the detector will trace out a well-defined track. This track must be analyzed to determine whether it could be realistic, given the spacecraft's position and speed. Rapid analysis of apparent motions that are consistent with real satellites may be a natural one to perform onboard the spacecraft with AI algorithms.

We expect any satellites, on first detection, to be essentially point-sources, unresolved by the camera. One will have information on the brightness of the objects -- cosmic ray hits may vary widely from one frame to the next, while satellite images should have relatively constant integrated signal levels. Extended objects will be easier to verify, but only if they can be seen in multiple images, because cosmic ray hits can sometimes cluster.

Preliminary Results

We have constructed a prototype onboard satellite search algorithm and have successfully tested it on the Ida-Dactyl images from Galileo. The onboard automated process detects a satellite in the presence of similar-appearing features such as: background stars, detector defects, and cosmic ray hits. At the farthest range the satellite is one pixel in size and only a few digital intensity levels higher than background, and lower in intensity than many of the cosmic ray hits. The detection is performed long before the encounter to encompass the entire Hill sphere, to avoid rapid changes in geometry due to spacecraft motion, and to allow time for possible processing and replanning. The small body is detectable by the process only if geometry allows for visual separation of the asteroid from the satellite body, and if the satellite is not obscured by the asteroid. The process relies on temporal information, therefore several rapidly sampled frames are necessary for proper execution. The satellite detection process is capable of incorporating short image sequences where images are taken in different filters. The process was tested on all the available images of Ida and Dactyl. Its performance was perfect. No manual parameter selection was necessary (all the parameters were selected autonomously by built-in procedures). The farthest sequence available was collected at Inc. distance of 171,318 km to the asteroid center, 3 hours 50 minutes before the closest approach. Initial estimates of false detection probabilities, given this set of images and circumstances, range from 0.001% to 0.07%, depending on the filter used for the observations.

AUTOMATED ANALYSIS AND DECISION MAKING FOR PLANETARY ULTRAVIOLET SPECTROSCOPY

Science Objectives

One of the most commonly used and powerful techniques for the exploration of planetary and cometary atmospheres is ultraviolet spectroscopy. UV spectrographs have been carried to every planet visited by spacecraft. UV spectra derived from these instruments are useful for obtaining the compositions and first-order physical properties (e.g., temperature, pressure) of upper planetary atmospheres and cometary comae.

UV spectra are also useful for categorizing the excitation mechanisms of planetary airglows, electrolights, aurora, and other phenomena, and for certain types of surface photometry.

We are developing a flight experiment to demonstrate automated first-order interpretation of UV spectra for compositional information. We intend to design and implement a decision-making algorithm that uses the output of this automated UV Spectral Interpretation Tool Experiment (UV-SITE) to select among several retargetable observation options during a mission flyby, to determine which would be scientifically most valuable, based on the data being obtained. Our goals are (i) to develop the algorithms required to make automated compositional inferences from UV spectra; (ii) to implement these algorithms with flight-like software onboard a flight-like microprocessor; and (iii) to demonstrate their functionality in a flight experiment, as an enabling capability for future NASA missions.

Scientific Benefits

From a scientific standpoint, the development of automated interpretation of UV spectra would allow the spacecraft to decide for itself which to pursue when there are many potential avenues of scientific value. For example, this capability could detect when the spectrum of a given region has sufficient signal-to-noise (so the aperture can be moved on to the next locale). As another example, this capability could allow near-real-time determination of whether faint features are present in the spectrum (so that the instrument executive can decide whether it is profitable to remain on this location (or spectral bandpass), or whether it would be wiser (i.e., of higher value) to explore other bandpasses or spatial locations. Although these kinds of capabilities would be useful on practically any mission involving atmospheric spectroscopy, it holds particular promise for encounters like the 1991 S-1 comet encounter or other flybys, where the reaction time from Earth is too slow to "close the loop" on unique mission science opportunities which perish after the flyby.

From an operations standpoint, the automated interpretation of UV spectra to determine their quality (i.e., S/N at various wavelengths) and content (i.e., the identification of the emissions being detected) would also enable the intelligent prioritization of datasets for transmission to the ground (or to varying degrees of compression prior to transmission). Since the scientific success of planetary missions is directly proportional to the value of the data returned, it is clear that this kind of analysis onboard the spacecraft has the potential to improve the usefulness of the many datastream-limited missions envisioned for the future.

Problem Description

The scientific problem of UV spectral classification is well-known. Simply put, the software needs to determine where in the spectrum statistically significant peaks occur (or in some cases the absence of candidate feature peaks on a prescribed list). The strength of the peak is a measure of the amount of emission. The ratios of diagnostic peaks (preselected by researchers on the ground), once identified by their position in the spectrum, yield measures of the relative abundances of different atmospheric constituents (or in the case of ratios among features of the same species, information about the excitation mechanisms acting in and the thermal properties of the atmosphere) being studied.

By comparing the results obtained during different instrument pointings, it is possible to "map out" the vertical or horizontal structure of emissions in an atmosphere, thus allowing mice to choose interesting regions to return to, or to predict the locations of transition zones between regions which are useful for study.

Technical Approach

The specific algorithm we propose for the UV-SITE autonomy experiment is easy to describe. Simply put, we will remove flat field and cosmic rays effects, establish an instrumental wavelength scale, and then convert the resulting count (or count rate) spectra to absolute units using the instrument transfer function (i.e., a stored sensitivity curve). Using these "reduced data" one then searches a statistically significant presence or absence of emission features in a table specific to the observation target. Features that are present (or notably absent) are thus catalogued. Using the catalog of features and the ratios of preselected,

diagnostic features, the spectrum can be evaluated to determine its content. From that content, either deterministic or fuzzy-logic based rules can be used to evaluate the worth of the spectrum for transmission to Earth and what the spectrum says about the relative value of retargetable observations. For example, the discovery of noble gases in a cometary coma would be so important that it would likely outrank additional observations of tail emissions or more mapping of H₂O content in the coma.

Given the increasing complexity of new-generation, micro UV spectrometers, it is also important to point out that instrument reconfiguration to choose between various wavelength regions, or to make a choice between spectral and spatial coverage is becoming increasingly common. However, most scientists agree that the choices one might make in hindsight (i.e., after or during a flyby) are often different from preconceived notions. As a result, a tool that allows the instrument to be reconfigured in response to the data it is obtaining would open a whole new world of adaptive scientific exploration at a distance.

Once the UV spectral recognition and scoring tool is validated in the flight environment, we expect future missions to take advantage of it to make decisions on retargetable observations, how much to compress (or whether to even transmit) specific images stored in mass memory, and how to adapt the instrument configuration during an encounter or orbital mission, thereby enhancing the scientific return and lowering mission operations costs.

Preliminary Results

This experiment is less mature at this time; we will report on preliminary results in a future forum.

FUTURE PROSPECTS

New Exploration and the Planetary Web

Our technological vision for meeting the challenge of returning quality science through limited communications bandwidth and at low cost will actually put scientists in a more direct link with the spacecraft than has been possible in the past. Technologies like pattern recognition and machine learning will place a part of the scientist's awareness onboard the spacecraft at launch time through trainable object recognizers. Such awareness can be used to prioritize downlink or to autonomously time-critical follow-up observations without ground interaction. Onboard data mining methods can be used to include candidate discoveries in each downlink for scientists' scrutiny. When a scientist determines that a discovery has been made, a new recognizer can be quickly developed and uploaded, and the reprioritized mission continues, modified by the scientist in a much more intimate and efficient manner than is possible today.

We may expect autonomy for science to evolve through several types of mission scenarios. Beginning with the agile, solitary explorer, this scenario would typically involve the intelligent selection of spaceborne scientific experiments by rapid analysis of preliminary low-resolution data. Feature detection algorithms can rapidly extract important objects such as volcanoes and craters (on planetary surfaces, possible satellites in orbit around asteroids, or hints of particular elements in UV spectra from cometary tails. Rapid decisions by automatic software can allow autonomous spacecraft responses to maximize scientific return, e.g., higher resolution imaging of areas rich in volcanoes and craters, redirection of spacecraft to confirm satellite presence and to map its orbit, or instructions to direct the UV spectrometer to collect higher-resolution data in interesting regions of a cometary tail. This capability will be especially important for fly-by missions or for studying transient phenomena.

Future mission scenarios are likely to involve small constellations of homogeneous platforms conducting a single mission, such as a free-flying interferometer composed of several elements. One goal for such a mission is the detection of neighboring planetary systems. Machine learning and data mining techniques would enable the detection of these systems with an order of magnitude higher sensitivity than would be possible with standard techniques. Future scientific analysis requires the development of techniques to analyze interferometric signals in the presence of uncertainty, e.g., ranking hypotheses by incorporation of raw signal with prior models of possible planetary systems. This enables both the characterization of planetary orbits and estimates of masses for planetary system candidates.

Finally, we can imagine a heterogeneous collection of space and ground platforms with varying capabilities, such as a proposed "Mars Society". A typical scenario here involves rapid response to important serendipitous scientific events. Imagine a fleet of spacecraft in orbit around a planet with a range of detecting devices. One detector images a volcanic region, and detects a possible eruption in the visible regime. The spacecraft coordinates with another spacecraft in the fleet carrying an I-R imager, which retargets and checks the same area, and confirms or discounts the event. If confirmed, several other craft in the fleet are redirected to focus on this important event to maximize science return. Ground vehicles take seismic and other data and may attempt to observe the phenomenon at close hand.

In a final and continuing phase of autonomy development and insertion, technologies such as the ones discussed here will have evolved by 2010 to the point where human awareness is routinely projected to remote spaceborne platforms and surface vehicles. Such capabilities will ultimately enable new types of exploration missions. There will be a new class of vigilant missions, where spacecraft are endowed with the ability to recognize certain classes of science events and to report on them autonomously. There will also be a new class of long duration missions, where an investigation can take place over a decade or more at low cost. Examples of these new missions are a submarine explorer under Europa's ice crust which remelts to the surface and relays back what it has found, a Uranus orbiter which conducts studies of seasonal cycles which last more than 20 years, micro rovers/insects which burrow under the surface of Mars looking for evidence of life or surveying water content, again resurfacing to transmit data, a **ACILITY** lander which tries to determine the origin of water ice in the permanently shaded south pole crater, aerobots in the Titan atmosphere where communication may be a challenge, and ultimately, spacecraft sent entirely out of the solar system. From an exploration perspective, autonomy is an extremely exciting prospect. As another logical consequence of such expanded capability, autonomous space platforms and surface vehicles (both planetary- and Earth (1)-based) as intelligent agents will represent much more sophisticated versions of today's World-Wide Web nodes. Direct and transparent access to the data returned from these widespread agents will be available to scientists, school children and the general public. Thus autonomy technology may also be expected to play a central role as a fundamental and rich source of information in our future's infocentric society.

Acknowledgements

The research described in this paper was carried out in part by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- Burl, M.C., Fayyad, U.M., Perona, P., Smyth, P., and Burl, M.P. (1994) "Automating the hunt for volcanoes on Venus" in *Proc. Computer Vision and Pattern Recognition Conference, CVPR-94*.
- Djorgovski, S., Weir, N., and Fayyad, U. (1994) "Processing and Analysis of the Palomar S^TScI Digital Sky Survey Using a Novel Software Technology", in D. Crabtree, J. Barnes, and R. Hanisch (eds.), *Astronomical Data Analysis Software and Systems III*, A.S.I.--CoSet.-Sect. 52, 39.
- Durda, D.D. (1996). "The formation of asteroidal satellites in catastrophic collisions." *Icarus*, 120, 217-219.
- Fayyad, U., N. Weir, and S.G. Djorgovski (1993). "SKICAT: a machine learning system for automated cataloging of large scale sky surveys." *Proc. of Tenth Int. Conf. on Machine Learning*, Morgan Kaufman.
- Fiesq, L., et al (1996). *Proceedings of the 19th AAS Guidance & Control Conference, Advances in the Astronautical Sciences*, R.D. Culp, ed., Breckenridge, CO.

Gehrels, T., Drummond, J. D., and Levenson, N. A. (1987). "The absence of satellites of asteroids." *Icarus* 70, 257-263.

Head, J. W., et al. (1991). "Venus volcanic centers and their environmental settings: recent data from Magellan." *EOS* 72, p.175, American Geophysical Union Spring meeting abstracts.

Van Flandern, T. C., Tedesco, E. F., and Binzel, R. J. (1979). Satellites of asteroids. in "Asteroids", ed. T. Gehrels (Tucson: Univ. of Arizona Press), pp. 443-465.

Weidenschilling, S., Paolicchi, P., and Zappala, V. (1989). Do asteroids have satellites? In "Asteroids II", eds. R. P. Binzel, T. Gehrels, and M. S. Matthews (Tucson: Univ. of Arizona Press), pp. 643-658.